

EXHIBIT 122

Creating a Competitive Advantage for AdX through Clean Ad Traffic

ADSPAM CONFIDENTIAL - PLEASE DO NOT SHARE WITHOUT PERMISSION

Status: (Draft - WIP)

Authors: dejager@, vegardj@

Last Updated: 5 May, 2014

The Problem: AdX is being unfairly held back and is losing to AppNexus

In 2013 AppNexus was selling 4–5x more impressions and winning advertiser spend at a 7x faster rate than Non-AdWords AdX

According to competitor analysis by JP Park and PDS combined with our own data, the two leading display ad exchanges, AdX and AppNexus, are on very different growth trajectories.

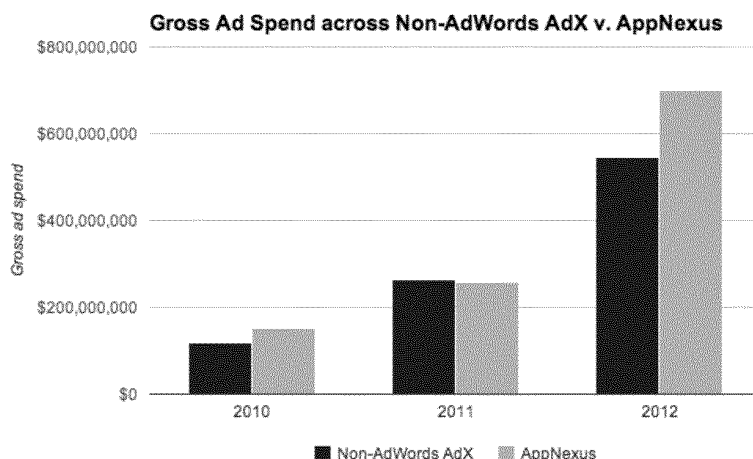
In 2010 gross spend across AppNexus was 25% larger than gross spend across Non-AdWords AdX: respectively, \$152 million versus \$120 million.

Between 2010 and 2011 gross spend across Non-AdWords AdX was growing twice as quickly as gross spend across AppNexus. In 2011 gross spend across Non-AdWords AdX drew level with gross spend across AppNexus, at just over \$250 million.

In September, 2012, growth in spend across AdX was reported in the PDS analysis to be slowing. By contrast, according to the revenue prediction for 2012 provided by AppNexus, spend across AppNexus was accelerating. Given the slowing growth in AdX spend, the revenue prediction provided by AppNexus was treated with skepticism: “AppNexus is also feeling the pinch. Investors are not buying their growth story.”

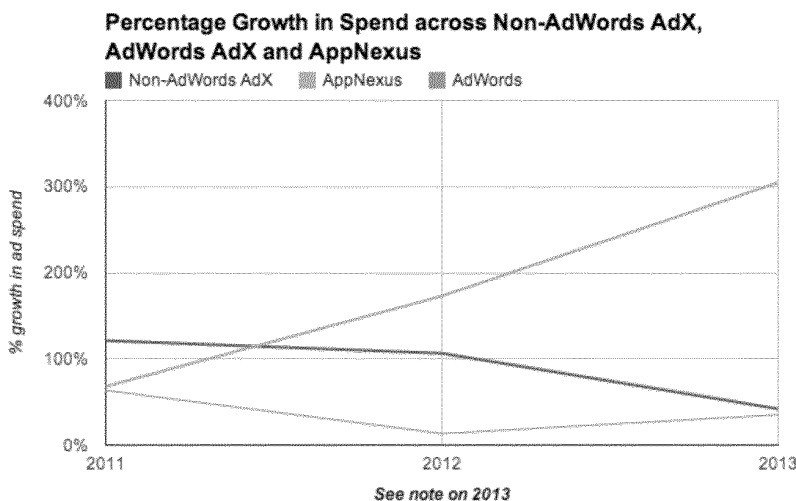
In January, 2013, AppNexus reported gross spend of \$700 million over 2012. This reported gross spend was in line with the revenue prediction of \$128,333,333 provided by AppNexus during the PDS analysis—with predicted revenue being 18.3% of reported gross spend.

Gross spend across Non-AdWords AdX over 2012 was \$550 million.



In July, 2013, Non-AdWords AdX was selling 2.6 billion impressions/day with gross spend across Non-AdWords AdX averaging \$2.1 million per day [AdX dashboard]. At the same time AppNexus was selling 13 billion impressions/day¹ (based on conversations between spider.io and AppNexus) with gross spend across AppNexus estimated to be \$7.8 million per day.²

The plot below shows the changing growth in spend across each of Non-AdWords AdX and AppNexus. The plot also shows the changing growth in spend across AdWords AdX. To provide a like-for-like comparison the growth in spend for 2013 across each of Non-AdWords AdX, AppNexus and AdWords AdX is measured as the daily gross spend in July, 2013, relative to the daily gross spend over 2012.



¹ NB Thirteen billion ad impressions were being sold per day through the AppNexus exchange and forty billion ad impressions were being served per day by AppNexus infrastructure.

² We used three data points to estimate the average CPM of media sold by AppNexus: (i) the average CPM of media bought by DBM from AppNexus, (ii) the average CPM of exchange notifications provided by AppNexus to DBM and (iii) the average CPM for media sold by AppNexus as reported by SiteScout.

Higher inventory quality is not being rewarded with more advertiser spend

The growing difference in spend across Non-AdWords AdX versus AppNexus is surprising.

Higher quality ad inventory means a more predictable return on marketing investment for advertisers. This in turn should yield an increase in advertiser demand. (c.f. [market for lemons](#))

In particular, non-viewable ads or ads served to bots deliver no real value to advertisers. These are waste impressions, which make return on marketing investment more unpredictable for advertisers. So if a display ad exchange reduces the risk of an advertiser buying these types of waste impression, then this should yield an increase in advertiser demand.

This is surprisingly not happening in practice.

AdX is not being rewarded with more advertiser spend relative to AppNexus. This is despite (1) AdX offering advertisers the opportunity to buy ad impressions with [Active View CPM](#) as the billable unit (thereby removing uncertainty over ad viewability) and (2) AdX also being regarded across the industry as providing significantly more protection against fraudulent and illegitimate inventory than any other exchange—in particular, relative to AppNexus.

[The most recent list of Active View CPM buyers](#) is illustrative. The leading direct-response display buyers—Rocket Fuel, Quantcast, DBM and DataXu—are conspicuous in their absence from this list. The list does include leading brand-focused display ad networks like Collective, Undertone and Cox Digital. However, only 0.67% of spend and only 0.72% of impressions across the whole of AdX were based on Active View CPM in April, 2014 [[PowerDrill details](#)].

It is perhaps worth adding a qualifying note about fraudulent and illegitimate ad traffic. Whilst we believe that AdX currently provides more protection against fraudulent and illegitimate inventory than AppNexus, we do not know this. We have anecdotal evidence to suggest that AppNexus sells significant volumes of spam inventory. We know also that AppNexus has very few engineers focused on inventory quality. However, no broad quantitative analysis has been performed into the quality of AppNexus inventory—for example, across the AppNexus inventory bought by DBM. To boot, whilst we are working toward deploying sophisticated protections against spam impressions across AdX—in particular, against malware-generated spam impressions—only derivative and basic protections are currently in place. AdSpam currently identifies and filters spam *clicks* across AdX. This allows AdSpam to impute spam classifications for impressions with the same cookie as some spam click. AdSpam also has blacklisting and various policy-based filters in place across AdX. If some impression or some conversion has no associated click event, and there is no conflicting blacklist fact or policy-based filter, then this impression or conversion will fall outside of AdSpam's current analysis. This is clearly problematic as impressions are the billable unit across AdX and value is typically attributed to media buyers on a post-view basis—which does not require a click event.

Perceived campaign performance does not reflect *real* campaign performance

We have explored this counterintuitive behaviour by display media buyers—and, in particular, by the direct-response buyers who continue to dominate the demand side.

Before the acquisition of spider.io we worked closely with the senior management of companies like Rocket Fuel, Quantcast, and DataXu to determine how best to incorporate viewability measurement and fraud detection into the optimisation engines that dictate their buying strategies.

The senior management of these companies believe that viewability measurement and fraud detection will be strategically important in the long term. However, all have run internal tests to show that in the short term incorporating viewability measurement into their optimisation engines would actually lead to these companies losing money. They have also made clear that they do not want to incorporate fraud detection in the short term as they expect this would similarly lead to their losing money.

The reason for these surprising verdicts about the short term is that display advertisers reward their direct-response media buyers for hitting CPA targets—typically post-view CPA, but also post-click CPA. This is problematic because media buyers can unfortunately exploit both non-viewable impressions and bot impressions to inflate the CPA numbers reported in DCM, Adometry or Convertro dashboards. These are the dashboards that advertisers currently use to measure and compare the campaign performance achieved by different media buyers.

For demand-side ad networks like Rocket Fuel and Quantcast any change in reported CPA directly impacts profit. These companies are paid on a CPA basis—their business model being CPM to CPA arbitrage. For demand-side platforms like DBM and DataXu any change in reported CPA indirectly impacts profit. These companies typically charge a fixed proportion of media spend—rather than selling media on a CPA basis. However, advertisers hold demand-side platforms like DataXu accountable to CPA targets. Failure to hit these targets results in customer attrition.

The problem of gamed CPA metrics is well known in the industry. Retargeters are particularly maligned for cookie bombing across low-viewability media to intercept credit for conversions rather than trying to drive more conversions. Companies like Quantcast, in particular, are vocal about this unscrupulous behaviour. But retargeters are not the only problem. In quiet corners many direct-response media buyers confess that they turn a blind eye when their optimisation engines buy media below Zynga games or at the bottom of eBay auction pages. This type of media often has viewability rates of less than 1%, however it provides a particularly cost-effective way to drop cookies and intercept credit for conversions.

Bot traffic is similarly problematic for the CPA numbers reported in the DCM, Adometry or Convertro dashboards used by advertisers. DCM, for example, does very little spam filtering when determining CPA numbers. Even the most basic filter to prevent too many clicks from the same cookie is not yet in place—meaning that DCM is less protected against spam than the old

DFA6. (c.f. Code Limerick.) This means that bot traffic often looks very good according to the numbers reported in today's dashboards.

Bot traffic is not just a problem for direct-response media buyers. In analysis across one of the largest video ad exchanges, 38% of the ad inventory was being served to bots. The buyers of this inventory are almost entirely brand buyers and their reach and frequency metrics will have been spammed by the bots.

Proposed Solution: Ensure that perceived campaign performance more closely reflects real campaign performance

Filter waste impressions from the attribution models of DCM, DBM and Google Analytics

If perceived campaign performance tended towards real campaign performance, then advertiser spend should tend towards higher quality ad inventory. In particular, if the CPA dashboards used by advertisers to measure campaign performance only attributed value to ad impressions that are both (i) served to legitimate human visitors and (ii) viewable, then advertiser spend should start to move to higher quality ad inventory.

A group of us are currently working toward filtering spam and non-viewable impressions from DCM's attribution models. This is significant because most advertisers use DCM as their ad server—particularly in the US—and most of these advertisers use their DCM dashboard to measure campaign performance and to allocate budgets to media buyers.

This same group is also aiming to filter waste impressions from the attribution models of DBM and Google Analytics. In the case of DBM this is because customers are apparently increasingly using only their DBM dashboards when measuring campaign performance—without appealing to their DCM dashboards.

Lead public narrative on how spam and non-viewable impressions skew attribution

The increasing use of DBM dashboards to measure campaign performance alludes to a broader problem that still needs to be addressed. Attribution companies like Adometry, Convertro and C3 Metrics are trying to disintermediate DCM from the advertiser's decisions about budget allocation. Much like DBM, media buyers like Rocket Fuel and Quantcast are also trying to influence the advertiser's decisions on budget allocation by providing alternative attribution dashboards.

If changes are made to the attribution models of DCM, then these will need to be explained publicly. We will need to lead a public narrative on how spam and non-viewable impressions skew attribution. This will help reverse any disintermediation of DCM from the advertiser's decisions on budget allocation. It will also help prevent customer escalations when CPA metrics in DCM start to diverge from the CPA metrics shown in other dashboards—and, indeed, from the CPA metrics previously shown in DCM.

Lead public debate on tackling fraud—in particular, malware-driven fraud

The success of any public narrative on how spam and non-viewable impressions skew attribution will depend on how credible Google's stance is with respect to spam.

Historically Google's treatment of spam has been highly regarded within the industry, but there are indications that this can now no longer be taken for granted—as noted in [A Troubling Update from the TOGI Solutions Subgroup](#). There are now several third-party traffic auditing companies and these companies are incentivised to find fault with Google's handling of spam. There has also been a shift in the public conversation from spam, broadly, to botnet traffic, specifically. Botnets are the dominant problem across display. This is a conversation that we can and should be leading. If we do not contribute to the public discussion it will not be clear that our approach to hunting and defending against ad-fraud botnets is best in class.

Recommended Actions

AdSpam and Active View to work with the attribution teams of DCM, DBM and Google Analytics to expedite filtering of spam and non-viewable impressions from the respective attribution models

We recommend expanding and formalising [this working group](#) to expedite the group's efforts.

AdSpam to work with AdX and DBM to expedite integration of BotGuard across all AdX and DBM impressions

We are currently working with the DCM team toward integrating BotGuard across all DCM impressions. This will allow us to filter spam impressions from DCM's attribution models. We recommend expediting similar collaboration with the AdX and DBM teams to integrate BotGuard across all AdX and DBM impressions.

AdX to investigate other potential blockers of advertiser spend across AdX

If AdX does not currently offer significantly cleaner inventory than AppNexus, then AdSpam can enable this imminently. In collaboration with DCM and the Active View teams, AdSpam can also provide AdX with a competitive advantage through cleaner inventory. Implementing these changes will not necessarily lead to an increase in spend across AdX. There may be other reasons which also limit advertiser spend across AdX. We recommend that AdX explore all the potential blockers of advertiser spend.

AdSpam to create a public blog

We recommend that AdSpam creates a public blog along the lines of [Facebook's Protect the Graph blog](#), which focuses on:

- the mechanics of ad-fraud malware (e.g. the resurrection of ZeroAccess);
- the impact of ad-fraud malware on the display ad ecosystem (e.g. a quantitative analysis of how retargeting strategies are currently being gamed); and
- the impact of spam and non-viewable impressions on today's attribution models (including quantitative studies).

It has been suggested that Sasha Tsvyashchenko and Douglas de Jager write the blog posts.

We recommend that AdSpam recruits two analysts to work full-time on providing supporting research for the blog: a malware researcher and a quantitative analyst. For some context, this is what a co-founder of OkCupid had to say about blogging at OkCupid³: "The posts each took 4-8 weeks of full-time work for [Christian Rudder] to write. Plus another 2-4 weeks of dedicated programming time from someone else on the team. It's easy to look at an OkTrends post, with all its simple graphs and casual writing style and think someone just threw it together, but it probably had 50 serious revisions. And we threw out a lot of research that didn't turn into good posts."

We recommend that AdSpam establishes lines of support from Google PR and legal to maximise the impact of each post whilst mitigating any potential legal risks.

Scott Spencer and Vegard Johnsen to put forward new TOGI proposals

We recommend that Scott joins Vegard on the TOGI Solutions Subgroup and that they put forward new proposals to set the subgroup on a better path.

Vegard Johnsen to join Scott Spencer in speaking at display advertising conferences

We recommend that AdSpam produces an extensive PR Q&A document and that Vegard Johnsen, AdSpam's PM for display, uses this to lead the public narrative on malware and attribution at display advertising conferences.

³ OkCupid's blog, OkTrends, remains arguably the most famous startup blog. The blog was read by millions, with each post averaging 32,500 Facebook likes and 4,222 tweets. This is an example post: <http://goo.gl/UFqQmP>.